

Exploiting natural dynamics in robot control

Matthew M. Williamson*

MIT AI Lab, 545 Technology Square Rm 828, Cambridge, MA 02139

<http://www.ai.mit.edu/people/matt>

email: matt@ai.mit.edu

Abstract

This paper presents an approach to robot arm control based on exploiting the dynamical properties of an adaptive oscillator circuit coupled to the joints of an arm. The approach is implemented on a real robot arm, and swings pendulums at their natural frequencies, turns cranks and manipulates slinky toys. These actions are all achieved using the same architecture, without any modeling of the arm or its environment. The simple nature of the oscillators, and the lack of modeling results in a robust and very simple system.

1 Introduction

This paper presents an approach to the control of robot arms which exploits the physical coupling of the arm and its environment. Dynamical characteristics of the arm are exploited rather than modeled, allowing a very simple control scheme to exhibit a variety of interesting rhythmic behaviors. The system is implemented on the arms of the humanoid robot Cog [Brooks and Stein, 1994], which is illustrated in Figure 1.

Robot arms are complex systems, with non-linear kinematics and dynamics. They interact with an environment that can also be dynamically complex, making the whole system awkward to model and control. In spite of the complexity, there are likely to be portions of the dynamics that are particularly suited for certain tasks. These parts may be configurations where the dexterity of the arm is increased (e.g., away from singularities), or areas where motions are inherently smooth, require low energy or low control effort. Exploiting these parts of the robot dynamics should result in a scheme which by working with rather than




Figure 1: Cog playing with a slinky toy. This picture shows the humanoid form of the robot, with the two 6 DOF arms used in this paper. The robot is using its elbow joints to move the slinky, exploiting the physical structure of the slinky to coordinate the two arms (see section 5).

against the natural dynamics of the system, reaps the benefits of simplicity, robustness, and efficiency.

To design such a scheme, there are a number of options; one can model the system accurately and use optimization techniques to find the best ways to move (the approach of [Uno *et al.*, 1989]), one can explore the system and remember good solutions, or one can design a system that exploits the natural dynamics directly [Greene, 1982]. A further refinement is the ability to shape the dynamics of the whole system, so that “good” solutions are always used.

The system described in this paper exploits directly the natural dynamics of the arm and its environment, and shapes the dynamics of the system to be appropriate to the task at hand. Key to the approach are the properties of an adaptive oscillator circuit, that interacts with the dynamics of the arm giving useful

*This work is funded in part by the Office of Naval Research. The author is supported by an EPSRC studentship.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 1998		2. REPORT TYPE		3. DATES COVERED 00-00-1998 to 00-00-1998	
4. TITLE AND SUBTITLE Exploiting natural dynamics in robot control				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street The Strata Center, Building 32, Cambridge, MA, 02139				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

and robust behavior. Also important is the spring-like behavior of the actuators used at the joints. The oscillators are used to drive the joints of the arm, using proprioceptive information (position and force) to adapt their behavior. As the arm moves in the environment, loads are transmitted through the structure of the arm allowing the oscillators along the whole length of the arm to respond. By changing the type of proprioception used, the overall behavior of the arm can be altered.

The system has been used to drive pendulums at their natural frequencies, turn cranks and play with slinky toys using the same architecture, without any explicit modeling or calculation.

The rest of the paper describes in detail the implemented system, giving a description of the behavior in a number of different situations. This behavior is then analyzed in more detail, highlighting how the dynamics of the arm are exploited, as well as providing biological evidence for the chosen solution.

2 The arms and the oscillators

The two 6 degree of freedom arms used in this work are mounted on the humanoid robot Cog [Brooks and Stein, 1994]. The arms have been specially designed to interact stably and robustly with unstructured environments. Each joint actuator is force controlled, and consists of a series elastic actuator [Pratt and Williamson, 1995]. These provide low noise force control, shock tolerance and a provably stable interaction with passive environments [Colgate and Hogan, 1989].

At the joints of the arm, a simple proportional-derivative position control loop is used, making the desired torque at the i th joint

$$\tau_i = k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i \quad (1)$$

where k_i is the stiffness of the joint, b_i the damping, θ_i the joint angle and θ_{vi} the equilibrium point. The dynamical characteristics of the arm can be changed by altering the stiffness and damping of the arm, and the posture of the arm can be changed by altering the equilibrium points [Williamson, 1996]. This type of control preserves stability of motion, and since the inner torque control is provided by the series elastic actuators, the overall system is both compliant and shock resistant, making it easy to operate the arm in unstructured environments.

The oscillator model consists of two simulated neurons arranged in mutual inhibition. The model for the neuron is taken from [Matsuoka, 1985], and describes the firing rate of a real biological neuron with self-inhibition:

$$\begin{aligned} \tau_1 \dot{x}_1 &= -x_1 - \beta v_1 - \omega y_2 - \sum_{j=1}^{j=n} h_j [g_j]^+ + c \quad (2) \\ \tau_2 \dot{v}_1 &= -v_1 + y_1 \quad (3) \end{aligned}$$

$$\tau_1 \dot{x}_2 = -x_2 - \beta v_2 - \omega y_1 - \sum_{j=1}^{j=n} h_j [g_j]^- + c \quad (4)$$

$$\tau_2 \dot{v}_2 = -v_2 + y_2 \quad (5)$$

$$y_i = [x_i]^+ = \max(x_i, 0) \quad (6)$$

$$y_{out} = y_1 - y_2 \quad (7)$$

where x_i is the firing rate, v_i is a variable representing the self-inhibition of the neuron (modulated by the adaption constant β), and the mutual inhibition is controlled by the parameter ω ¹. The output of each neuron y_i is taken as the positive part of x_i , and the output of the whole oscillator is y_{out} . The input g_j is arranged to excite one neuron and inhibit the other, by applying the positive part ($[g_j]^+$) to one neuron and the negative part to the other. The inputs are scaled by gains h_j . The other important parameters are c , a constant that specifies the amplitude of the output, and two time constants τ_1 and τ_2 . For steady oscillations, τ_1/τ_2 should be in the range 0.1–0.5, the oscillator having an output frequency $w_{osc} \propto 1/\tau_1$. A detailed analysis of this type of oscillator was published by [Matsuoka, 1985; 1987].

The oscillator is connected to the robot joints by using the output y_{out} to move the equilibrium point θ_v . The oscillations are about a fixed posture θ_p , making the equilibrium point

$$\theta_v = y_1 - y_2 + \theta_p = y_{out} + \theta_p \quad (8)$$

For the examples in this paper, the inputs to the oscillators are taken to be either the force (τ) or the position (θ) of the joint, as shown in Figure 2. These signals in general have an offset (due to gravity loading, or oscillation about a non-zero posture), so when the positive and negative parts are extracted to be applied to the oscillators, a high pass filter is used to remove the dc component.




Figure 2: System schematic. The oscillator input is either the torque at the joint τ_i or the angle of the joint θ_i , giving two different modes of operation. The output is used to drive the equilibrium position θ_{vi} of the joint, and so move the arm. There is no “set-point” for the complete system, the system behavior coming from the interaction between the oscillator and arm dynamics.

¹Values used $\omega = 2.5$ and $\beta = 2.5$

Condition	Position Feedback	Torque Feedback
Free	Drives at resonant frequency of the overall system	Drives at low frequency
Position constrained	Resists at low frequencies, drives at high frequencies	Equilibrium point tracks motion (low impedance)
Force perturbation	Ignores	Entraines and drives limb approximately in phase with perturbation

Table 1: Summary of oscillator behavior

The behaviour of the complete arm-oscillator system (Figure 2) is complex and is examined in more detail in [Williamson, 1998]. There is no “set-point” for the control; the final behavior of the system emerges from the interaction between the oscillator and the arm dynamics, without any adaption of parameter values. Under different feedback conditions, and different dynamical situations of the arm, the overall system exhibits different behaviors. Table 1 summarizes the behaviour of the overall system under two different types of feedback, and three conditions for the arm: free to move, constrained to move, and under a force perturbation.

When the actuated limb is free to move, under position feedback the oscillator drives the system at its resonant frequency, while under torque feedback, the system is driven at a fairly constant low frequency. This occurs for a wide range of system natural frequencies for a single set of oscillator parameters.

When the limb is forced to move in a rhythmic manner, the oscillator senses the motion and responds to the perturbation. Over a wide range of frequencies (typically in the range $0.1w_{osc}$ to $4w_{osc}$) the oscillator tracks the frequency of the perturbation. Under position feedback, at low frequencies the oscillator actively resists the imposed motion, while at high frequencies it actively helps the motion. Under torque feedback, the oscillator causes the equilibrium point to track the external motion, resulting in a very low torque. This low impedance behavior is observed below w_{osc} , above which the oscillator cannot track the perturbation, and shuts off.

When a force perturbation is applied to the limb (such as from some external source, like a hand moving the arm), under position feedback, the force is ignored. However under force feedback, for frequencies in the range $0.1w_{osc}$ to w_{osc} the oscillator entrains the frequency of the perturbation, and moves the limb to be approximately in phase with the disturbance. Above the natural frequency of the oscillator, the disturbance is no longer entrained.

The following sections describe in more detail the behavior of the oscillators and the arm, and show how they exploit the dynamical properties of the arm to

give stable behavior without any computation.

3 Rhythmic motions

Human arms (and Cogs arms) can be thought of as masses connected by springs, whose frequency response makes the energy and control required to move the arm vary with frequency. At the resonant frequency, the control need only inject a small amount of energy to maintain the vibration of the mass of the arm segment on the spring of the muscles and tendons. The frequency response of the system thus determines speeds and frequencies that efficiently move the arm, which can then be exploited giving simple, efficient control.

It appears that humans exploit the natural frequencies of their systems, swinging pendulums at “comfortable” frequencies equal to the natural frequency [Hatsopoulos and Warren, 1996].




Figure 3: Performance of the robot swinging a pendulum. The graph shows the frequency that the oscillator drove the pendulum plotted against the natural frequency of the pendulum. The action of the oscillator to drive the pendulum at its natural frequency over the range 5 to 9 rad/s. The natural frequency of the oscillator (w_{osc} - horizontal solid line) is 7 rad/s making the entrainment range about 60%

One could find the best frequency to drive the system by using conventional system identification tech-

niques, and then driving at a fixed frequency [Ljung, 1987], however the behavior of the neural oscillators can achieve the same result without explicit computation. Figure 3 shows the result of a single robot joint driving a pendulum with different lengths and joint stiffnesses². The figure shows that over the range 5 to 9 rad/s for the particular case of $w_{osc} = 7$ rad/s the behavior of the joint is to tune into the natural frequency of the system.

The oscillator behavior is robust, driving a variety of different systems at their individual natural frequencies for a single set of parameters. When the robot is swinging the pendulum it is resistant to perturbations, returning quickly to the correct frequency. Since the system works because of the tight coupling between the dynamics of the oscillator and the actuated system, the response to changes in either system is very quick. In addition the computation required is very small. From a wider perspective, the behavior of the oscillators of automatically finding the most efficient speed for the motion will be useful in designing controllers which exploit the natural dynamics of the system.

This work is similar to the approach of [Hatsopoulos *et al.*, 1992], although the oscillator system described in this paper has been extended to more applications.

4 Arm Dynamics

The dynamical loads and forces experienced by even a simple arm are extremely complicated [Hollerbach and Flash, 1982], and for the biological system, with its non-linear musculature, the picture is even more complex. To cope with this complexity, robots are often designed to be either very lightweight [Salisbury *et al.*, 1988], use complex dynamical models [An *et al.*, 1988], or are simply moved slowly. The system presented here exploits the interaction forces to communicate between the joints of the robot, giving coordinated multi-dof motion without any kinematic modeling.

Humans also exploit the interaction forces, either by working with them to get smooth, approximately linear motion [Flash and Hogan, 1985], perhaps by using the cerebellum [Uno *et al.*, 1989; Kawato, 1993]. During development [Thelen *et al.*, 1992] and during learning of a motor task [Schneider *et al.*, 1989], humans learn to compensate for, and work with the interaction forces. Humans also exploit the physical structure of the skeleton to reduce the forces at their joints when carrying heavy loads.

The system presented in this paper exploits the physical structure of the arm and the interaction forces to coordinate motion along the full length of

the arm. The oscillators only receive local proprioception, so can only affect one another through the interaction forces in the limb. Since the oscillators can entrain motion and forces over a range of frequencies, the joints can be actuated at different frequencies, but become coordinated through the arm. This is illustrated in Figure 4, showing the behavior of the arm while flailing. When the proprioception is turned on, the joints are coordinated, while when it is off, they move at different frequencies. This result shows the power of the oscillator behavior coupled with the arm dynamics to give stable coordinated motion.




Figure 4: Flailing. Both plots show the angle of the shoulder (solid) and the elbow (dashed) as the speed of the shoulder is changed (speed parameter dash-dot). The top graph shows the response of the arm without proprioception, and the bottom with proprioception. The synchronization is clear in the lower graph, the only connections between the joints being through the physical structure of the arm.

Exploiting this communication, the arm has also been used for multi-dof constrained tasks, such as crank turning and pumping a bicycle pump. Crank turning can be achieved by driving one joint of the arm, the constraint of the crank being communicated through the physical structure of the arm to force the motion of the other joints. Depending on the feedback used (see Table 1), those joints can either track the motion, or actively drive it, both resulting in the crank being turned. The oscillator response is only at the entrained frequency, and for the entrained motion, causing perturbations to be rejected, and giving stability in the task space. The motion of the arm is along a low impedance trajectory, requiring little energy to make the motion. In this sense the oscillators shape the dynamics of the arm to make them “good” in the task space.

The crank turning does work, but is limited with respect to the size and the position of the crank. Robotic

²For a pendulum actuated through a spring the natural frequency is given by $w = \sqrt{g/l + k/I}$, g being gravity, l the length of the pendulum and I its inertia.

crank turning solutions (using hybrid position/force control [Raibert and Craig, 1981] or impedance control [Hogan, 1985]) can turn different cranks in different locations, but require kinematic modeling and are thus sensitive to modeling errors. By solving the problem in a limited domain, we are showing the possibility of exploiting the characteristics of the oscillators and the physical situation to perform the task without any kinematic knowledge about the arm³. Whether this approach can be extended to different cranks in different positions is a question for further work.

This approach to multi-dof freedom motion is similar to that of [Taga *et al.*, 1991], who made a simulated biped walk using a set of neural oscillators which entrained the dynamics of the legs.

5 Dynamics of objects

A general purpose arm needs to interact with objects in the world, requiring the controller to not only exploit the dynamical characteristics of the arm, but also those of manipulated objects. These have mass and inertia (pendulums, baseball bats etc), spring-like properties (drums, spongy materials), or simply their own dynamics (bouncing balls, juggling clubs, etc). Humans are remarkably adept at interacting with and perceiving the dynamical properties of objects [Turvey and Carello, 1995].

As with arm dynamics, a modeling technique can be used although it is difficult to generalize [Mason and Lynch, 1993]. An alternative is to explore the dynamics of the individual systems, as suggested by [Schaal and Atkeson, 1993]. He described “open loop stable” schemes for a variety of juggling tasks, which are schemes where perturbations to the system are corrected by the natural dynamics, without requiring any reactive corrections. Closed loop control built around this open loop stability was found to be simple and robust.

Since the oscillators receive information through the arm of the robot, they cannot distinguish whether loads are from the arm or from objects. Their ability to entrain perturbations over a range of frequencies also makes them useful for interacting with dynamical objects. A stable behavior has been found in using two arms to play with a slinky toy, as shown in Figure 1.

When the toy is passed from hand to hand the weight of the toy causes the force on the hands to change. The oscillators can track this force perturbation, and move the hands in phase with it, so moving the slinky. Figure 5 shows the drive to the two hands with and without the force feedback, showing that the




Figure 5: The robot operating the slinky. Both plots show the outputs from the oscillators as the proprioception is turned on and off. When it is on, the outputs are synchronized, and when off, the oscillators move out of phase. The only connection between the oscillators is through the physical structure of the slinky.

motion of the slinky is enough to very quickly lock the phase and the frequency of the two oscillators.

Since stable operation of the slinky compliments the behavior of the oscillators, the overall performance is stable to perturbations. In addition, since the synchronization is through the physical mass of the slinky, the behavior is robust to the use of different joints and different configurations of the arm. By exploiting the dynamics of the slinky, a very simple and robust scheme is obtained.

6 Summary

This paper has presented an number of examples of the behavior of a simple oscillator circuit driving the joints of a compliant robot arm. In each case, the coherent behavior displayed is a direct consequence of the physical coupling between the oscillators, the arm and the environment. The oscillator senses and responds to the resonant frequency of the system when driving the pendulum, it uses the dynamical forces along the limb to coordinate the joints for multi-dof motions, and uses the dynamics of the slinky toy to synchronize the playing behavior. In addition to the coupling, the properties of the oscillator, and the spring-like properties of the arm contribute to the robust and flexible performance of the system. The lack of kinematic modeling required makes the system very computationally simple.

This behavior is generated by a simple two neuron oscillator using only position and force feedback, begging the question of what might be possible using more neurons, connections between the neurons,

³Kinematic knowledge about the crank is provided, albeit not in an explicit form, by putting the crank in the hand of the robot, and deciding the approximate amplitudes of the joint motions

or maybe different networks of neurons that can be switched in and out. Different proprioceptive signals from touch sensing, vision, or even audition could also be used. How a more complex scheme could be designed, adapted or even evolved to exploit the rich dynamics of the system is a topic for further work. By using a more complex architecture, but still exploiting the dynamics of the system, a wider range of interesting behaviors will hopefully be easy to achieve.

References

- [An *et al.*, 1988] Chae H. An, Christopher G. Atkeson, and John M. Hollerbach. *Model-based control of a robot manipulator*. MIT Press, Cambridge, MA, 1988.
- [Brooks and Stein, 1994] Rodney A. Brooks and Lynn A. Stein. Building brains for bodies. *Autonomous Robots*, 1(1):7–25, 1994.
- [Colgate and Hogan, 1989] Edward Colgate and Neville Hogan. An analysis of contact instability in terms of passive equivalents. In *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*, pages 404–409, 1989.
- [Flash and Hogan, 1985] Tamar Flash and Neville Hogan. The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7):1688–1703, July 1985.
- [Greene, 1982] Peter H. Greene. Why is it easy to control your arms? *Journal of Motor Behavior*, 14(4):260–286, 1982.
- [Hatsopoulos and Warren, 1996] Nicholas G. Hatsopoulos and William H. Warren. Resonance tuning in rhythmic arm movements. *Journal of Motor Behavior*, 28(1):3–14, 1996.
- [Hatsopoulos *et al.*, 1992] Nicholas G. Hatsopoulos, William H. Warren, and J. N. Sanes. A neural pattern generator that tunes into the physical dynamics of the limb system. In *Proceedings of the International Joint Conference on Neural Networks*, volume 1, pages 104–109, Baltimore, MD, 1992.
- [Hogan, 1985] Neville Hogan. Impedance control: An approach to manipulation. *Journal of Dynamic Systems, Measurement, and Control*, 107:1–24, 1985.
- [Hollerbach and Flash, 1982] John M. Hollerbach and Tamar Flash. Dynamic interactions between limb segments during planar arm movement. *Biological Cybernetics*, 44:67–77, 1982.
- [Kawato, 1993] Mitsuo Kawato. Inverse dynamics model in the cerebellum. In *Proc Intl Joint Conf on Neural Networks*, pages 1329–1335, Baltimore, MD, 1993.
- [Ljung, 1987] Lennart Ljung. *System identification : theory for the user*. Prentice-Hall, 1987.
- [Mason and Lynch, 1993] Matthew T. Mason and Kevin M. Lynch. Dynamic manipulation. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems (IROS-93)*, pages 152–159, Yokohama, Japan, 1993.
- [Matsuoka, 1985] K. Matsuoka. Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biological Cybernetics*, 52:367–376, 1985.
- [Matsuoka, 1987] K. Matsuoka. Mechanisms of frequency and pattern control in neural rhythm generators. *Biological Cybernetics*, 56:345–353, 1987.
- [Pratt and Williamson, 1995] Gill A. Pratt and Matthew M. Williamson. Series elastic actuators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-95)*, volume 1, pages 399–406, Pittsburg, PA, July 1995.
- [Raibert and Craig, 1981] Marc H. Raibert and J. J. Craig. Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, 103:126–133, 1981.
- [Salisbury *et al.*, 1988] J. Kenneth Salisbury, William T. Townsend, Brian S. Eberman, and D. M. DiPietro. Preliminary design of a whole arm manipulation system (wams). In *Proc 1988 IEEE Intl Conf on Robotics and Automation*, 1988.
- [Schaal and Atkeson, 1993] Stefan Schaal and Christopher G. Atkeson. Open loop stable control strategies for robot juggling. In *Proceedings 1993 IEEE International Conference on Robotics and Automation*, volume 3, pages 913–918, 1993.
- [Schneider *et al.*, 1989] Klaus Schneider, Ronald F. Zernicke, Richard A. Schmidt, and Timothy J. Hart. Changes in limb dynamics during the practice of rapid arm movements. *Journal of Biomechanics*, 22(89):805–817, 1989.
- [Taga *et al.*, 1991] G. Taga, Y. Yamaguchi, and H. Shimizu. Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological Cybernetics*, 65:147–159, 1991.
- [Thelen *et al.*, 1992] Esther Thelen, Ronald Zernicke, Klaus Schneider, Jody Jensen, Kathy Kamm, and Daniela Corbetta. The role of intersegmental dynamics in infant neuromotor development. In G. E. Stelmach and J. Requin, editors, *Tutorials in Motor Behavior II*, pages 533–548. Elsevier Science Publishers, 1992.
- [Turvey and Carello, 1995] Michael T. Turvey and Claudia Carello. Dynamic touch. In W. Epstein and S. Rogers, editors, *Perception of Space and Motion*, pages 401–490. Academic Press, 1995.
- [Uno *et al.*, 1989] Y. Uno, M. Kawato, and R. Suzuki. Formation and control of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, 61:89–101, 1989.
- [Williamson, 1996] Matthew Williamson. Postural primitives: Interactive behavior for a humanoid robot arm. In *Fourth International Conference on Simulation of Adaptive Behavior*, pages 124–131, Cape Cod, Massachusetts, 1996.
- [Williamson, 1998] Matthew M. Williamson. A simple approach to rhythmic robot control. 1998. Submitted to ICRA '98.